



Sound Event Detection in Multichannel Audio using Convolutional Time-Frequency-Channel Squeeze and Excitation

Wei Xia^{1,2}, Kazuhito Koishida²

¹ Center for Robust Speech Systems, University of Texas at Dallas, TX 75080

² Microsoft Corporation, One Microsoft Way, Redmond, WA 98052

wei.xia@utdallas.edu, kazukoi@microsoft.com

Abstract

In this study, we introduce a convolutional time-frequency-channel "Squeeze and Excitation" (tfc-SE) module to explicitly model inter-dependencies between the time-frequency domain and multiple channels. The tfc-SE module consists of two parts: tf-SE block and c-SE block which are designed to provide attention on time-frequency and channel domain, respectively, for adaptively recalibrating the input feature map. The proposed tfc-SE module, together with a popular Convolutional Recurrent Neural Network (CRNN) model, are evaluated on a multi-channel sound event detection task with overlapping audio sources: the training and test data are synthesized TUT Sound Events 2018 datasets, recorded with microphone arrays. We show that the tfc-SE module can be incorporated into the CRNN model at a small additional computational cost and bring significant improvements on sound event detection accuracy. We also perform detailed ablation studies by analyzing various factors that may influence the performance of the SE blocks. We show that with the best tfc-SE block, error rate (ER) decreases from 0.2538 to 0.2026, relative 20.17% reduction of ER, and 5.72% improvement of F1 score. The results indicate that the learned acoustic embeddings with the tfc-SE module efficiently strengthen time-frequency and channel-wise feature representations to improve the discriminative performance.

Index Terms: Sound event detection, squeeze and excitation, attention, multichannel audio, convolutional recurrent neural network

1. Introduction

Sound event detection (SED) task involves labeling the time stamps of a sound event in audio streams and detecting the sound type. Speech and non-speech sounds such as laughter and music contains lots of useful information. Being able to detect environmental sound events in multi-channel audios can greatly help us understand surrounding acoustic environments and enables many applications such as audio surveillance and rare sound detection [1, 2, 3]. It is also helpful to improve the performance of speech enhancement and separation systems if we could know the types of sounds [4, 5]. Robotic systems can employ SED for navigation and natural interaction with surrounding acoustic environments [6, 7]. Smart home devices can benefit from it for environmental sound understanding [8, 9]. Sound event detection is attracting more and more attention nowadays from these applications, as well as multimedia content retrieval [10, 11] and audio segmentation [12, 13, 14].

Real life audio recordings typically have many overlapping sound events, the task of recognizing all the overlapping sounds is considered as polyphonic SED. Lots of efforts have

been proposed to address this task by predicting frame-wise labels of each sound event class. Models like Gaussian Mixture Model (GMM) [15], Hidden Markov Model (HMM) [16], Recurrent Neural Networks (RNN) [17, 18], and Convolutional Neural Networks (CNN) [19, 20] have been explored extensively for this task. More recently successful results were obtained by stacking CNN, RNN and FC layers consecutively, referred jointly as the convolutional recurrent neural network (CRNN) [21]. In order to improve the recognition of overlapping sound events, several multi-channel SED methods have also been proposed. For example, Sharath [22, 23, 24] showed that Deep Neural Network can directly learn from low-level features such as generalized cross-correlation with phase based weighting (GCC-PHAT) for multi-channel sound detection.

The basic building block for most SED models is the convolutional layer, which learns filters capturing local spatial patterns along all the input channels and generates feature maps jointly encoding the time-frequency and channel information. A lot of work has aimed at improving the joint encoding of spatial and channel information [25], but much less attention has been given towards encoding of time-frequency and channel-wise patterns independently with domain information. Some recent work attempted to address this issue by explicitly modeling the inter-dependencies between the channels of feature maps. One promising approach to accomplish this is a component called "Squeeze & Excitation" (SE) block [26, 27], which can be seamlessly integrated into the CNN model. This SE block factors out the spatial dependency by global average pooling to learn a channel specific descriptor, which is used to rescale the input feature map to highlight only useful channels.

In this study, we introduce a time-frequency-channel Squeeze and Excitation (tfc-SE) block for multi-channel sound event detection. Different from the original SE block mentioned above, which uses global average pooling and only excites on the channel axis, we adapt and extend it to both the time-frequency and channel domains. Multi-channel audio contains different information at each time-frequency location, for example, we may pay more attention to high energy parts in the spectrogram. We first introduce a time-frequency SE (tf-SE) block that aims to adaptively recalibrate the learned feature map. It does not change the receptive field, but provides time-frequency attention to certain regions. Further, we propose two methods to combine the channel wise SE and time-frequency SE. These methods aggregate the unique properties of each block and make feature maps to be more informative on both domains. We show that with the best tfc-SE block, error rate of the SED system decreases from 0.2538 to 0.2026, relative 20.17% reduction of ER. It also has 5.72% relative improvement of F1 score compared to the original CRNN.

In the following sections, we describe the tfc-SE approach and corresponding baseline systems in Section 2. We provide

This work is done under Microsoft Research Internship program.

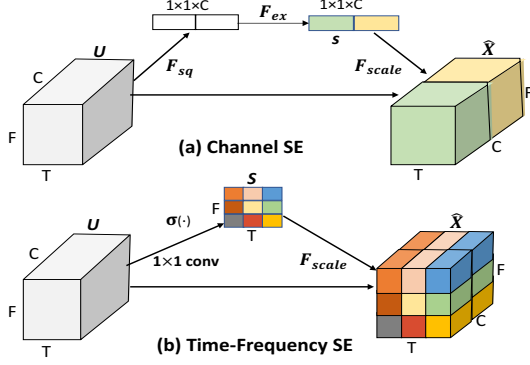


Figure 1: Channel SE and Time-Frequency SE blocks.

detailed explanations of our experiments in Section 3, as well as results and discussions in Section 4. Finally we conclude our work in Section 5.

2. Sound Event Detection Systems

2.1. Baseline Convolutional RNN

We use a recently proposed Convolutional Recurrent Neural Network (CRNN) [21] to learn the acoustic representations of multi-channel audio signals for sound event detection. The CRNN model has three components, the convolutional layers to learn the time-frequency representations of audio waveform, the recurrent layers to learn the temporal information, and the final classification layer for sound event detection. The model configuration is given in Table 1. We use three layers of 2D CNN to learn the shift invariant features from multi-channel spectrograms. Each CNN layer has P filters of a 3×3 kernel. After each CNN layer, we use a ReLU activation function and batch normalization to normalize the activations. A Maxpooling layer is applied on the frequency axis and the sequence length of T -frame input features are kept unchanged.

Table 1: Baseline CRNN model architecture, when input are $T = 256$ frames, 256 dimensional and 16 channels features, $P = 64$ filters, $Q = 128$ nodes, $R = 128$ nodes, and $N = 11$ sound types.

Type	Filter shape	Input shape
Conv1	$3 \times 3 \times 16 \times 64$	$256 \times 256 \times 16$
Maxpool1	1×8	$256 \times 256 \times 64$
Conv2	$3 \times 3 \times 16 \times 64$	$256 \times 32 \times 64$
Maxpool2	1×8	$256 \times 4 \times 64$
Conv3	$3 \times 3 \times 16 \times 64$	$256 \times 4 \times 64$
Maxpool3	1×2	$256 \times 2 \times 64$
Bi-GRU1	-	256×128
Bi-GRU2	-	256×128
FC1	128×128	256×128
FC2	128×11	256×128

The output activation from CNN is further fed to bidirectional RNN layers which are designed to learn the temporal context information from the CNN output activations. Specifically, Q nodes of Gated Recurrent Units (GRU) are used in each layer with tanh activations. For classification, we use two fully connected (FC) layers. The first FC layer has R nodes each with

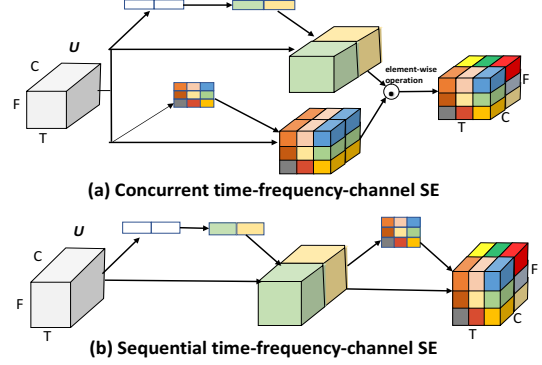


Figure 2: Concurrent time-frequency-channel SE and sequential time-frequency-channel SE blocks.

linear activation. The last FC layer consists of N nodes with sigmoid activations, each corresponding to one of the N sound event classes to be detected.

In the following sections, we introduce the ‘‘Squeeze & Excitation’’ (SE) blocks. We insert the SE block after each convolutional layer to adaptively recalibrate the feature representations.

2.2. Channel wise squeeze and excitation

2.2.1. Squeeze operation

The channel-wise SE block, c-SE, is illustrated in Fig. 1 (a). In order to model the inter-dependencies between multiple channels of audio signals, we firstly define a squeeze operation to embed the global time-frequency information into a channel descriptor. We consider the input feature map of the multi-channel audio as $U = [U_1, U_2, \dots, U_c]$ where $U_c \in \mathbb{R}^{T \times F}$ is the feature matrix of channel c . We use a global average pooling layer to generate a channel-wise vector $\mathbf{z} \in \mathbb{R}^{1 \times 1 \times C}$ with its c -th element,

$$z_c = \mathcal{F}_{sq}(U_c) = \frac{1}{T \times F} \sum_i \sum_j U_c(i, j) \quad (1)$$

This operation embeds the global time-frequency information into the vector \mathbf{z} . This vector contains statistics that are expressive of the whole time-frequency input.

2.2.2. Excitation operation

In order to capture channel wise dependencies, a gating mechanism with a sigmoid activation function is used to learn a non-linear relationship between channels,

$$\mathbf{s} = \mathcal{F}_{ex}(\mathbf{z}, \mathbf{W}) = \sigma(g(\mathbf{z}, \mathbf{W})) = \sigma(\mathbf{W}_2 \delta(\mathbf{W}_1 \mathbf{z})) \quad (2)$$

where $\mathbf{W}_1 \in \mathbb{R}^{\frac{C}{r} \times C}$ and $\mathbf{W}_2 \in \mathbb{R}^{C \times \frac{C}{r}}$ are weights of two fully-connected layers and δ is the ReLU operator. The dimensional reduction factor r indicates the bottleneck in the channel excitation block. Note that the original channel dimension is recovered by the second FC layer. With a sigmoid layer σ , the channel-wise attention vector \mathbf{s} is obtained. Finally we recalibrate U with the attention vector as,

$$\hat{X}_c = \mathcal{F}_{scale}(U_c, s_c) = s_c \cdot U_c \quad (3)$$

$\hat{X} = [\hat{X}_1, \hat{X}_2, \dots, \hat{X}_C]$ is the final channel wise recalibrated features. In this block, the input features are attentively scaled

so that important channels are emphasized and less important one are diminished.

2.3. Time-frequency squeeze and excitation

Here we introduce the time-frequency wise squeeze and excitation block (tf-SE), shown in Fig. 1 (b). The concept is similar to the c-SE and the main difference is to compress the feature map U along the channel axis using a 1-by-1 convolution [28]. The excitation operation is performed on the time-frequency map. We assume that the time frequency space may contain more information for SED.

We write the input feature map U in an alternate form,

$$U = \begin{bmatrix} \mathbf{u}^{1,1} & \mathbf{u}^{1,2} & \dots & \mathbf{u}^{1,F} \\ \mathbf{u}^{2,1} & \mathbf{u}^{2,2} & \dots & \mathbf{u}^{2,F} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{u}^{T,1} & \mathbf{u}^{T,2} & \dots & \mathbf{u}^{T,F} \end{bmatrix} \quad (4)$$

where $\mathbf{u}^{i,j} \in \mathbb{R}^{1 \times 1 \times C}$ is the feature bin at the time-frequency location (i, j) . The squeeze operation of the tf-SE is done using a 1-by-1 convolution which can be represented by the linear combination of all channels C at a location (i, j) ,

$$s^{i,j} = \sigma((\mathbf{w}^{i,j})^T \cdot \mathbf{u}^{i,j}) \quad (5)$$

where $\mathbf{w}^{i,j}$ is the filter coefficients and $s^{i,j}$ is the (i, j) -th element of the squeezed matrix $\mathbf{S} \in \mathbb{R}^{T \times F}$. We use a sigmoid function σ to limit the range of the matrix \mathbf{S} to $[0, 1]$, which is used to recalibrate the features on the time-frequency domain. Each value $s^{i,j}$ corresponds to the relative importance in the time-frequency space.

Similar to the c-SE, the excitation of the tf-SE is carried out as

$$\hat{\mathbf{x}}^{i,j} = \mathcal{F}_{scale}(\mathbf{u}^{i,j}, s^{i,j}) \quad (6)$$

and $\hat{\mathbf{x}}^{i,j}$ is the (i, j) -th element of the recalibrated output \hat{X} .

2.4. Concurrent and sequential time-frequency-channel SE

Each of the above explained c-SE and tf-SE blocks has its unique properties. The c-SE blocks recalibrates the channel information by incorporating global time-frequency information. On the other hand, the tf-SE block generates an time-frequency attention map, indicating where the network should focus more to aid the sound classification.

We propose two ways to combine the complementary information from these two SE blocks to form the time-frequency-channel SE (tfc-SE): 1) Concurrent recalibration, illustrated in Fig. 2 (a) and 2) Sequential recalibration in order of channel first and then time-frequency as in Fig. 2 (b).

For the concurrent tfc-SE, we present four ways to aggregate the c-SE and the tf-SE blocks.

- (i) Addition: the two recalibrated feature maps are added element wise with equal weights.
- (ii) Multiplication: the feature maps are multiplied element wise.
- (iii) Maximization: each location (i, j, c) of the output feature map has the maximum activation of the two feature maps element-wise.
- (iv) Concatenation: the two feature maps are concatenated along the channel axis, which means that the volume of the input feature is twice larger.

The concurrent tfc-SE and sequential tfc-SE with these aggregation strategies will be evaluated in the following sections.

3. Experimental Setup

3.1. Dataset

To study the effectiveness of the CRNN model with the tfc-SE module in multi-channel sound event detection, we used the synthetic eight-channel TUT Sound Events 2018 - Circular array, Reverberant and Synthetic Impulse Response (CRESIM) dataset [23]. The dataset synthesizes the DCASE 2016 task 2 dataset [29], which has 11 isolated sound event classes such as speech, door slam, phone ringing, coughing, and keyboard.

Each of the sound classes has 20 examples, of which 16 are randomly chosen for the training set and the rest 4 for the test set, in total 176 examples from 11 classes for training, and 44 for testing. We selected the O3 subset which has maximum three temporally overlapping sources. It consists of three cross-validation splits with 240 training and 60 testing recordings of length 30 sec sampled at 44100 Hz. In this dataset, circular microphone array recording is simulated with 8 omnidirectional microphones equally spaced on 5cm radius. More details can be found in [23].

3.2. Feature extraction

On frame basis, we compute the magnitude and phase of spectrograms and concatenate them along the channel axis as input features. A Hamming window of length M with 50% overlap is used to extract the spectrogram from each audio channel. The zeroth bin is excluded from the spectrogram, so that each frame produces a $M/2 \times 16$ feature matrix. The input feature is mean and variance normalized on the frame level. A sequence of the T -frame spectrograms are stacked and fed into the network.

3.3. Evaluation metrics

We evaluated our sound detection model using the standard polyphonic SED metrics, error rate (ER) and F-score calculated on segments of one second with no overlap as proposed in [30]. F1 (ideally F1 = 1) is based on true and false positives, and the error rate (ER) (ideally ER = 0) is based on the total number of active sound event classes in the ground truth. A joint SED error score can be considered as $S_{SED} = (ER + 1 - F)/2$.

3.4. Model configuration and training

We explored various CRNN models and feature configurations. It was found from preliminary experiments that the best input feature setting is a window length of $M = 512$ and sequence length of $T = 256$ (equivalent to 1.486 sec). The CRNN model was trained on the CRESIM dataset for 1000 epochs with a batch size of 64. Early stopping was applied if there is no improvement on the S_{SED} score after 100 epochs. We built a various models with different parameters such as the number of nodes of CNN/RNN/FC and CNN filter size. The best baseline setting can be found in Table 1. The last fully connected layer has 11 nodes with sigmoid activation and the cross entropy loss for classification was employed for network training. We used an Adam optimizer with the learning rate 10^{-3} . For each sound event class all the predicted probabilities in the sequence ($T = 256$) were examined at the frame level and the class detection was claimed as true if there is any probability which is larger than a threshold of 0.5.

4. Results and Discussions

4.1. Experimental results

In order to thoroughly evaluate our proposed methods, we will conduct detailed ablation analysis in this section. We first perform experiments on the c-SE, followed by the tf-SE, and then the sequential tfc-SE for sound event detection. We further investigate the aggregation strategies for the concurrent tfc-SE. Finally we analyze various factors that may affect the performance of SE blocks.

From Table 2, we observe that our proposed channel wise SE (c-SE) block improves both F1 and ER compared with the original CRNN model. After utilizing the c-SE block, overall error score S_{SED} of the model decreases from 0.2285 to 0.2194, relative 3.98% gain. This approach also achieves relative 6.30% improvement in ER. We also find a consistent performance improvement with the time-frequency SE (tf-SE) block, relative 4.08% and 12.02% improvement in F1 and ER, respectively, compared with the original CRNN model. The above results indicate that the tf-SE is more efficient than c-SE, which aligns with our assumption that the time-frequency space may have more meaningful information than channel for the SED task.

Finally, we test the tfc-SE block which combines the c-SE and t-SE activations. Both the concurrent and sequential models improve the performance of the original CRNN model by a large margin. The best result is obtained with the sequential tfc-SE block, which achieves 84.23% F1 score and 0.2026 error rate, relative 5.72% and 20.17% improvement compared with the original CRNN. In terms of the overall S_{SED} score, the sequential tfc-SE block outperforms the original CRNN model by 21.18% relatively.

Table 2: Averaged SED results on CRESIM overlap 3 three data splits using the CRNN model with time-frequency-channel squeeze and excitation blocks.

Model	F1 (%)	ER	S_{SED}
CRNN	79.67	0.2538	0.2285
+c-SE	79.91	0.2378	0.2194
+tf-SE	82.95	0.2233	0.1952
+tfc-SE concurrent	83.57	0.1982	0.1812
+tfc-SE sequential	84.23	0.2026	0.1801

4.2. Aggregation strategies

We further investigate the aggregation strategy of the concurrent tfc-SE block, among the four choices. We observe from Table 3 that all aggregation methods increase the SED performance against the original CRNN. Using maximization provides the best performance. As concatenation aggregation increases the model complexity by the doubled number of channels, maximization operator looks a better choice for a lower computa-

Table 3: SED results on CRESIM overlap 3 split 1 subset using CRNN + tfc-SE concurrent block with different aggregation operations.

Aggregation	F1(%)	ER	S_{SED}
Addition	84.92	0.1791	0.1649
Multiplication	84.48	0.1959	0.1756
Maximization	85.79	0.1703	0.1562
Concatenation	85.26	0.1841	0.1657

tional cost. For all other experiments, we use the maximization-based aggregation for tfc-SE blocks.

4.3. Sensitivity to dimension reduction ratio

The reduction ratio r introduced in Section 2.2.2 is an important hyperparameter that allows us to vary the capacity and computational cost of the c-SE blocks in the model. Table 4 reveals that the performance does not improve monotonically with increased capacity. This probably comes from overfitting due to a larger model capacity. We find that $r = 8$ is a good option and we use it for our other experiments.

Table 4: SED results on CRESIM overlap 3 split 1 subset using CRNN + tfc-SE concurrent block with different dimension reduction ratios.

Ratio r	F1(%)	ER	S_{SED}
2	85.07	0.1751	0.1622
4	85.59	0.1881	0.1661
8	85.79	0.1703	0.1562
16	83.51	0.1957	0.1803

4.4. Squeeze and excitation operator

For the squeeze operator, we examine the significance of using global average pooling as opposed to global max pooling (while keeping the excitation operator sigmoid) in Table 5. Though both max and average pooling are effective, average pooling achieves a little better result and we use it all in our paper. We next assess the excitation operator. Two options, ReLU and tanh, are experimented by replacing from the sigmoid (with leaving the squeeze operator global average pooling). It suggests that it is important to choose the sigmoid operator in order to make the tfc-SE block effective.

Table 5: SED results on CRESIM overlap 3 split 1 subset using CRNN + tfc-SE concurrent block with different squeeze and excitation operators.

Operator	F1(%)	ER	S_{SED}
Max	82.89	0.1971	0.1841
Avg	85.79	0.1703	0.1562
ReLU	83.42	0.1799	0.1728
Tanh	82.58	0.2196	0.1969
Sigmoid	85.79	0.1703	0.1562

4.5. Model complexity

The original CRNN model has 496,587 parameters. The tfc-SE block requires only 0.7% additional parameters, which is 500,067 in total.

5. Conclusions

In this paper, we proposed a convolutional time-frequency-channel squeeze and excitation block for multi-channel sound event detection, in order to model the feature inter-dependencies between channels and the time-frequency locations. The tfc-SE block was inserted after each convolution layer of the CRNN model. The proposed method was evaluated on the CRESIM dataset and it was shown to improve the original CRNN model in terms of both F1 score and error rate by a large margin. These results indicated that the tfc-SE block effectively recalibrates the feature maps adaptively by emphasizing more important channels and time-frequency locations.

6. References

- [1] J. Kotus, K. Lopatka, and A. Czyzewski, "Detection and localization of selected acoustic events in acoustic field for smart surveillance applications," *Multimedia Tools and Applications*, vol. 68, no. 1, pp. 5–21, 2014.
- [2] P. Foggia, N. Petkov, A. Saggese, N. Strisciuglio, and M. Vento, "Audio surveillance of roads: A system for detecting anomalous sounds," *IEEE transactions on intelligent transportation systems*, vol. 17, no. 1, pp. 279–288, 2016.
- [3] M. Crocco, M. Cristani, A. Trucco, and V. Murino, "Audio surveillance: A systematic review," *ACM Computing Surveys (CSUR)*, vol. 48, no. 4, p. 52, 2016.
- [4] D. Stowell, D. Giannoulis, E. Benetos, M. Lagrange, and M. D. Plumbley, "Detection and classification of acoustic scenes and events," *IEEE Transactions on Multimedia*, vol. 17, no. 10, pp. 1733–1746, 2015.
- [5] Q. Kong, Y. Xu, W. Wang, and M. D. Plumbley, "A joint separation-classification model for sound event detection of weakly labelled data," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 321–325.
- [6] R. Takeda and K. Komatani, "Sound source localization based on deep neural networks with directional activate function exploiting phase information," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2016, pp. 405–409.
- [7] W. He, P. Motlicek, and J.-M. Odobez, "Deep neural networks for multiple speaker detection and localization," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 74–79.
- [8] A. Southern, F. Stevens, and D. Murphy, "Sounding out smart cities: Auralization and soundscape monitoring for environmental sound design," *The Journal of the Acoustical Society of America*, vol. 141, no. 5, pp. 3880–3880, 2017.
- [9] C.-C. Kao, W. Wang, M. Sun, and C. Wang, "R-crmn: Region-based convolutional recurrent neural network for audio event detection," *arXiv preprint arXiv:1808.06627*, 2018.
- [10] M. Xu, C. Xu, L. Duan, J. S. Jin, and S. Luo, "Audio keywords generation for sports video analysis," *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 4, no. 2, p. 11, 2008.
- [11] Q. Jin, P. Schulam, S. Rawat, S. Burger, D. Ding, and F. Metze, "Event-based video retrieval using audio," in *Thirteenth Annual Conference of the International Speech Communication Association*, 2012.
- [12] A. Kumar and B. Raj, "Audio event detection using weakly labeled data," in *Proceedings of the 24th ACM international conference on Multimedia*. ACM, 2016, pp. 1038–1047.
- [13] M. Tian, G. Fazekas, D. A. Black, and M. Sandler, "On the use of the tempogram to describe audio content and its application to music structural segmentation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 419–423.
- [14] G. Wichern, J. Xue, H. Thornburg, B. Mechtley, and A. Spanias, "Segmentation, indexing, and retrieval for environmental and natural sounds," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 3, pp. 688–707, 2010.
- [15] P. K. Atrey, N. C. Maddage, and M. S. Kankanhalli, "Audio based event detection for multimedia surveillance," in *IEEE International Conference on Acoustics Speech and Signal Processing Proceedings (ICASSP)*, vol. 5. IEEE, 2006.
- [16] A. Mesaros, T. Heittola, A. Eronen, and T. Virtanen, "Acoustic event detection in real life recordings," in *European Signal Processing Conference*. IEEE, 2010, pp. 1267–1271.
- [17] T. Hayashi, S. Watanabe, T. Toda, T. Hori, J. Le Roux, K. Takeda, T. Hayashi, S. Watanabe, T. Toda, T. Hori *et al.*, "Duration-controlled lstm for polyphonic sound event detection," *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 25, no. 11, pp. 2059–2070, 2017.
- [18] G. Parascandolo, H. Huttunen, and T. Virtanen, "Recurrent neural networks for polyphonic sound event detection in real life recordings," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 6440–6444.
- [19] S. Hershey, S. Chaudhuri, D. P. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold *et al.*, "Cnn architectures for large-scale audio classification," in *IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 131–135.
- [20] H. Zhang, I. McLoughlin, and Y. Song, "Robust sound event recognition using convolutional neural networks," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 559–563.
- [21] E. Cakır, G. Parascandolo, T. Heittola, H. Huttunen, and T. Virtanen, "Convolutional recurrent neural networks for polyphonic sound event detection," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 6, pp. 1291–1303, 2017.
- [22] S. Adavanne, P. Pertilä, and T. Virtanen, "Sound event detection using spatial features and convolutional recurrent neural network," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 771–775.
- [23] S. Adavanne, A. Politis, J. Nikunen, and T. Virtanen, "Sound event localization and detection of overlapping sources using convolutional recurrent neural networks," *IEEE Journal of Selected Topics in Signal Processing*, 2018.
- [24] S. Adavanne, A. Politis, and T. Virtanen, "Multichannel sound event detection using 3d convolutional neural networks for learning inter-channel features," in *2018 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2018, pp. 1–7.
- [25] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei, "Deformable convolutional networks," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 764–773.
- [26] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [27] A. G. Roy, N. Navab, and C. Wachinger, "Concurrent spatial and channel squeeze & excitation in fully convolutional networks," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018, pp. 421–429.
- [28] M. Lin, Q. Chen, and S. Yan, "Network in network," *arXiv preprint arXiv:1312.4400*, 2013.
- [29] G. L. E. Benetos, M. Lagrange, "sound event detection in synthetic audio," [Online]. Available: https://archive.org/details/dcase2016_task2_train_dev, 2016.
- [30] A. Mesaros, T. Heittola, and T. Virtanen, "Metrics for polyphonic sound event detection," *Applied Sciences*, vol. 6, no. 6, p. 162, 2016.